

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

- 1 1. (Currently amended) A processor-implemented method for
2 enabling ~~efficient~~ communication between a first and a second node in a network
3 by routing network traffic through fault zones in the network, the method
4 comprising:
5 identifying a path from the first network node to the second network node,
6 wherein the path is a possible routing path for communication between the first
7 and second network nodes;
8 identifying a set of fault zones through which the identified path passes;
9 for each fault zone in the set of fault zones, assigning as a zone weight, the
10 number of paths from the first network node to the second network node that
11 include said fault zone, wherein the zone weight at each fault zone is an indicator
12 of the number of possible routing paths that will traverse that fault zone, and
13 wherein the zone weight enables routing decisions independent of network traffic
14 ~~is determined from the path configuration of the network;~~
15 calculating a path weight for the identified path, wherein said path weight
16 is equal to the sums of said zone weights for each fault zone included in the
17 identified path; and
18 selecting the identified path as the current routing path for routing the
19 network traffic from the first node to the second node,
20 whereby ~~efficient~~ communication from the first node to the second node is
21 enabled along the selected path.

1 2. (Original) The method of claim 1, further comprising:
2 identifying a new path from the first network node to the second network
3 node;
4 assigning zone weights to each fault zone in the new path;
5 calculating a new path weight for the new path; and
6 if the new path weight is lower than said path weight for the identified
7 path, selecting the new path as the current path for network traffic from the first
8 node to the second node.

1 3. (Previously presented) The method of claim 1, wherein:
2 the first network node is identified by a first identifier;
3 the second network node is identified by multiple identifiers, including a
4 second identifier;
5 selecting the identified path as the current path for network traffic from the
6 first node to the second node comprises selecting the identified path as the current
7 path for network traffic from the first node to the second node, wherein the second
8 node is identified by the second identifier; and
9 paths other than the identified path are selected as the current paths for
10 network traffic from the first node to the second node, wherein the second node is
11 identified by multiple identifiers other than the second identifier.

1 4. (Currently amended) A computer readable medium storing
2 instructions that, when executed by a computer, cause the computer to perform a
3 method for enabling ~~efficient~~ communication between a first and a second node in
4 a network by routing network traffic through fault zones in the network, the
5 method comprising:
6 identifying a path from the first network node to the second network node,
7 wherein the path is a possible routing path for communication between the first

8 and second network nodes;
9 identifying a set of fault zones through which the identified path leads;
10 for each fault zone in the set of fault zones, assigning as a zone weight, the
11 number of paths from the first network node to the second network node that
12 include said fault zone, wherein the zone weight at each fault zone is an indicator
13 of the number of possible routes that will traverse that fault zone, and wherein the
14 zone weight enables routing decisions independent of network traffic
15 ~~is determined from the path configuration of the network;~~
16 calculating a path weight for the identified path, wherein said path weight
17 is equal to the sums of said zone weights for each fault zone included in the
18 identified path; and
19 selecting the identified path as the current routing path for routing the
20 network traffic from the first node to the second node,
21 whereby ~~efficient~~ communication from the first node to the second node is
22 enabled along the selected path.

1 5. (Currently amended) A processor-implemented method for
2 enabling ~~efficient~~ communication between a first and a second node in a subnet by
3 determining routing between nodes in the subnet, the method comprising:
4 identifying multiple fault zones in the subnet, each fault zone comprising
5 one or more components of the subnet;
6 configuring a central subnet manager to manage routing between nodes in
7 the subnet;
8 identifying a set of paths from a first node having a first identifier to a
9 second node having multiple identifiers, including a second identifier, wherein the
10 set of paths are possible routing paths for communication from the first node to
11 the second node using any of the multiple identifiers, wherein traffic is deliverable
12 to the second node using any of the multiple identifiers;

13 for each fault zone traversed by one or more of the paths, establishing a
14 zone weight based on the number of paths from the first node to the second node
15 that traverse said fault zone, wherein the zone weight at each fault zone is an
16 indicator of the number of possible routing paths that will traverse that fault zone,
17 and wherein the zone weight enables routing decisions independent of network
18 traffic is determined from the path configuration of the subnet;
19 for each path in the set of paths, establishing a path weight from the sums
20 of the zone weights for each fault zone traversed by said path; and
21 for each of the multiple identifiers of the second node, selecting as the
22 current routing path from the first node to the second node identified by ~~each of~~
23 the said multiple ~~identifiers~~ identifier, from said set of paths, the path having the
24 best-least path weight,
25 whereby ~~efficient~~ communication from the first node to the second node is
26 enabled along the selected path.

1 6. (Currently amended) A computer readable medium storing
2 instructions that, when executed by a computer, cause the computer to perform a
3 method for enabling ~~efficient~~ communication between a first and a second node in
4 a subnet by determining routing between nodes in the subnet, the method
5 comprising:
6 identifying multiple fault zones in the subnet, each fault zone comprising
7 one or more components of the subnet;
8 configuring a central subnet manager to manage routing between nodes in
9 the subnet;
10 identifying a set of paths from the first node having a first identifier to the
11 second node having multiple identifiers, including a second identifier, wherein the
12 set of paths are possible routing paths for communication from the first node to
13 the second node using any of the multiple identifiers ~~wherein traffic is deliverable~~

14 ~~to the second node using any of the multiple identifiers;~~
15 for each fault zone traversed by one or more of the paths, establishing a
16 zone weight based on the number of paths from the first node to the second node
17 that traverse said fault zone, wherein the zone weight at each fault zone is an
18 indicator of the number of possible routing paths that will traverse that fault zone,
19 and wherein the zone weight enables routing decisions independent of network
20 traffic is determined from the path configuration of the subnet;
21 for each path in the set of paths, establishing a path weight from the sums
22 of the zone weights for each fault zone traversed by said path; and
23 for each of the multiple identifiers of the second node, selecting as the
24 current routing path from the first node to the second node, wherein the second
25 node is identified by ~~each of the said multiple identifiers~~ identifier, from said set
26 of paths, the path having the ~~best~~ least path weight,
27 whereby ~~efficient~~ communication from the first node to the second node is
28 enabled along the selected path.

1 7-13 (Canceled).

1 14. (Currently amended) A system for enabling ~~efficient~~
2 communication between nodes in a subnet by determining routing through the
3 subnet comprising multiple fault zones, comprising:
4 a network node configured to interface a client computing device with the
5 subnet, wherein each node is identifiable by one or more node identifiers; and
6 a client computing device, comprising:
7 a processor; and
8 a subnet manager module configured to determine routing between
9 a first node and a second node in the subnet, wherein the first node is
10 addressable by a first identifier and the second node is addressable by

11 multiple identifiers, including a second identifier;
12 wherein said subnet manager determines routing between the first node
13 and second node by:
14 for each fault zone in the subnet traversed by a possible routing
15 path from the first node to the second node, calculating a zone weight
16 based on the number of possible routing paths from the first node to the
17 second node that traverse said fault zone, wherein the zone weight at each
18 fault zone is an indicator of the number of possible routing paths that will
19 traverse that fault zone, and wherein the zone weight enables routing
20 decisions independent of network traffic~~is determined from the path~~
21 ~~configuration of the subnet;~~
22 for each of the paths from the first node to the second node,
23 calculating a path weight based on the sums of said zone weights for the
24 fault zones traversed by said path; and
25 selecting as the current routing path from the first node to the second node
26 identified by the second identifier, the path ~~r~~-having the best-least path weight,
27 whereby ~~efficient~~ communication from the first node to the second node is
28 enabled along the selected path.

1 15. (Original) The system of claim 14, wherein the client computing
2 device further comprises:
3 a memory configured to store path weights of current paths between
4 multiple pairs of node identifiers.

1 16. (Original) The system of claim 14, wherein said memory is further
2 configured to store, in association with each of the current paths, zone weights for
3 fault zones traversed by the current path.

1 17. (Original) The system of claim 14, wherein said subnet manager is
2 further configured to disseminate routing information to a plurality of nodes in the
3 subnet, said routing information including said current path from the first
4 identifier to the second identifier.

1 18-24 (Canceled).